ASCeNews Quarterly Newsletter
NA-ASC-500-09 Issue 13 SAND 2009-8321 W
December 2009

The Meisner Minute

Editorial by Bob Meisner

As we approach the end of 2009 and approach a lull in Washington budget activities, it seems an appropriate time to reflect on tri-lab achievements.

Over the past year, platform activities continued to provide new challenges and ground-breaking results. Roadrunner completed its shakedown period; not only is it ready to transition to classified operations, nine science projects produced results that are being published in technical journals. Dawn, the early delivery system for Sequoia, quickly hosted users and has been opened to the tri-lab. The facility housing the Red Storm machine was accredited as a SCIF and began hosting "non-traditional," national security simulations. Hyperion, the first supercomputer test-bed became operational, offering national collaborators unprecedented access to a large parallel resource for research and development work. The Capability Computing Campaign resource, Cielo that is scheduled to replace Purple in 2010, is being prepared for Critical Decision Level 2/3 approval (a federal process to begin system delivery). Finally, the Common Computing Environment is building synergy through a portfolio of projects that focus on providing a common software stack and configuration management for tri-lab use.

We have continued to develop and validate the ASC codes. The year 2009 saw the introduction of many new capabilities, including improved modeling of high-energy-density aboveground experiments in the Lawrence Livermore National laboratory (LLNL) nuclear design code system, the consolidation and extension of capabilities in the SNL Sierra framework, and the release of Los Alamos National Laboratory (LANL) Crestone project codes with capabilities to run Implicit Monte Carlo calculations on Roadrunner. We also demonstrated improvements in Uncertainty Quantification (UQ) capabilities; developing new algorithms, as well as providing initial demonstrations of the LANL Boost Validation Suite and LLNL Secondary Computational Assessment Metric Project.

Despite budget uncertainties at the beginning of 2009, the ASC execs tackled several challenges through unique tri-lab strategy projects. We also worked through a long continuing resolution period that lasted until March and induced uncertainty in our ability to achieve planned milestones. Yet, ASC completed all planned L1 and L2 milestones, with the exception of two which were canceled due to significant budget changes. Also completed was the NNSA stretch target and multi-site ASC milestone to deliver models and databases applicable to forensics and other nuclear security applications.

In last quarter's newsletter, I introduced the Exascale Initiative (EI) geared to field an exascale system to address DOE application codes by 2018. The Exascale Initiative Steering Committee (EISC) was asked to refine the case for exascale computing and produce a high level technology roadmap. Accordingly, the EISC held several workshops where computer vendors, industry, national laboratories, and academia explored the architectural and technology challenges that must be tackled to meet the needs of applications requiring exascale computing. Recently, Michael Strayer from the Office of Science and I presented a briefing on the EI to DOE undersecretaries, Steve Koonin and Tom D'Agostino. The presentation was well received and we are currently working on the steps required to make this a successful initiative.

All-in-all a banner year for the program because of your achievements. Your contributions keep the program vital and exciting. Thanks for another successful year and enjoy a well earned holiday season with friends and family.

New Research on Accelerating Simulations Demonstrates Notable Speedups

General-purpose graphics processing units (GPGPUs) have become popular as inexpensive accelerators for parallel computation and are considered important in the push towards exascale computing due to their low power requirements. As part of a research effort led by Mark Rintoul, Sandia scientists Mike Brown and Steve Plimpton implemented several algorithms for scalable GPGPU accelerated molecular dynamics (MD) simulation within the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) simulator.

Their work focused on acceleration for simulations using the Gay-Berne potential, a force-field designed for simulation of aspherical mesogens or coarse-grained particles. The scientists were able to demonstrate speedups of over 320x the single core CPU equivalent on the Tri-Lab Capacity Clusters (TLCC) and over 180x a single CPU simulation on Thunderbird. The 4-GPU simulations, requiring less than 1kW on a single node, were faster than 256-core simulations on either platform. The 256-core simulations require approximately 11.2 kW on a TLCC platform or 44.8 kW on Thunderbird. The improvements become even greater when the number of particles or interaction cutoff is increased.

The approach uses the existing spatial decomposition in LAMMPS to allow scaling to multiple CUDA-enabled GPUs on multiple nodes. The computation is divided further with a force-decomposition that assigns calculation of pair-wise particle interactions to each core on the GPU. This approach might be advantageous for improving strong scaling in MD; with one algorithm, the researchers found that greater than 20x speedups could be achieved with less than 1 particle per GPU core.

GPGPU acceleration for the Gay-Berne and Lennard-Jones potentials is currently available in LAMMPS. With funding support from Oak Ridge National Laboratory's (ORNL) Center for Accelerated Application Readiness, Sandia scientist Paul Crozier is leading a multi-institutional (Sandia, ORNL, Nvidia, Cray, Rensselaer Polytechnic Institute, Temple University, Georgia Tech) effort to expand and improve GPGPU capabilities in LAMMPS during FY10.

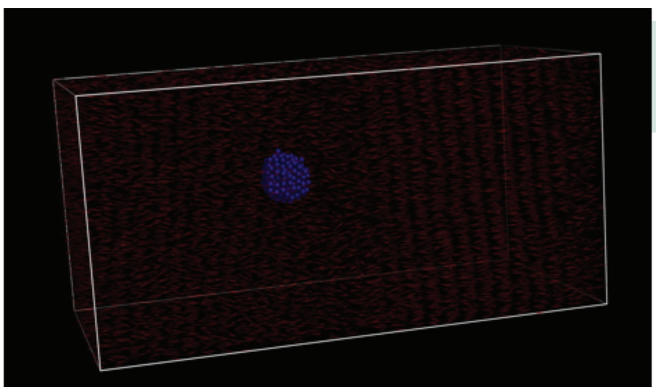

Snapshot from a GPGPU-accelerated simulation of a colloidal particle immersed in a nematic solvent.

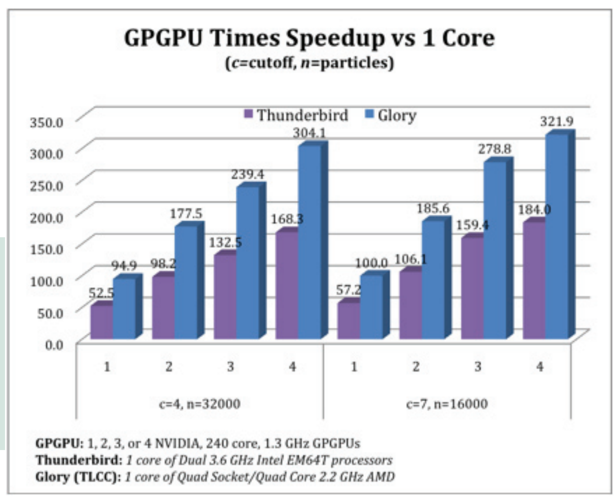

GPGPU speedup data from simulations of ellipsoidal mesogens using the Gay-Berne potential (microcanonical ensemble). GPGPU accelerated simulations are run on a single node with 4 GPUs. Interactions are only computed for particles closer than the cutoff distance, c.

$$U = U_r(\mathbf{A}_1, \mathbf{A}_2, \mathbf{r}_{12})\,\eta_{12}(\mathbf{A}_1, \mathbf{A}_2)\,\chi_{12}(\mathbf{A}_1, \mathbf{A}_2, \hat{r}_{12})$$

$$U_r = 4\varepsilon\left[\left(\frac{\sigma}{h_{12}+\gamma\sigma}\right)^{12} - \left(\frac{\sigma}{h_{12}+\gamma\sigma}\right)^{6}\right]$$

$$\eta_{12} = \left[\frac{2 s_1 s_2}{\det\left[\mathbf{A}_1^T \mathbf{S}_1^2 \mathbf{A}_1 + \mathbf{A}_2^T \mathbf{S}_2^2 \mathbf{A}_2\right]}\right]^{\upsilon/2}$$

$$s = \left[a_i b_i + c_i c_i\right]\left[a_i b_i\right]^{1/2}$$

$$\chi_{12} = \left[2\hat{\mathbf{r}}_{12}^T\left(\mathbf{A}_1^T \mathbf{E}_1 \mathbf{A}_1 + \mathbf{A}_2^T \mathbf{E}_2 \mathbf{A}_2\right)^{-1}\hat{\mathbf{r}}_{12}\right]^{\mu}$$

Gay-Berne potential for dissimilar biaxial ellipsoids. In the equation, **S** *represents the shape matrix,* **A** *represents the rotation matrix,* **E** *characterizes the relative well depths, and* $\mathbf{h}_{12}$ *represents a calculation to determine the distance of closest approach between the ellipsoids.*

For more information about this Roadrunner open science project, go to "Saturation of Backward Stimulated Scattering of Laser in the Collisional Regime" at http://www.lanl.gov/orgs/hpc/roadrunner/pdfs/openscience/Abstracts3%201.pdf.

Los Alamos' Roadrunner Open Science a Grand Success

Roadrunner at Los Alamos National Laboratory is the world's first petaflop hybrid supercomputer and number 4 on the Nov. 2009 TOP Green500 List of the world's most energy-efficient supercomputers.

In the spring and summer of 2009, Los Alamos accomplished the open-science aspect of application code efforts to prepare the Roadrunner advanced architecture machine for classified computing to begin in December 2009. Of the ten open-science projects using eight codes, seven of the projects were clear successes. Significant scientific accomplishments and speedups were achieved.

Roadrunner open science is gaining attention from dissemination of the results in Los Alamos press releases, conference presentations, and journal publications. Los Alamos is organizing a mini-symposium at the SIAM Conference on Parallel Processing and Scientific Computing (PP10), February 24–26 in Seattle, WA. Also, as announced at SC09, a Hybrid Multicore Consortium has been formed by Los Alamos, Oak Ridge, and Lawrence Berkeley national laboratories.

More information about the Roadrunner, can be found at the following site: http://www.lanl.gov/orgs/hpc/roadrunner/index.shtml.

Taking Charge of the "Impossible" at Supercomputing Conference (SCO9)

Taking on a charged particle physics calculation previously labeled "impossible," a team of Lawrence Livermore scientists developed an unorthodox strategy to exploit the power of massively parallel supercomputers and break new ground in scientific simulation. The multidisciplinary team, which included an IBM researcher, presented its methodology at SC09 and were finalists for the prestigious Gordon Bell Prize.

As supercomputing moves into the petascale (quadrillions of operations per second) era, scientists face the growing challenge of how to effectively use the increasing number of central processing unit cores to run more detailed simulations of scientific phenomena over longer time scales. In a reversal of the conventional practice of dividing up a problem and distributing it equally across the machine, Lab scientists carved up the problem according to the varied computational requirements needed to scale the individual component algorithms of the simulation. New BlueGene/P node technology, allowed them to use this new approach—called heterogeneous decomposition—to more fully exploit the system's capabilities.

This new capability was developed on two IBM BlueGene/P systems: the 500-teraFLOPS Sequoia Initial Delivery (ID) System (Dawn) at Lawrence Livermore and the 1.03-petaFLOPS JUGENE at Germany's Julich Supercomputing Center.

The team took on a problem that has long challenged scientists—a full understanding of the interaction of highly correlated charged particles. Until the team's recent simulations, molecular dynamics simulations involving electrostatic interactions were of insufficient length and time scale to fill the gaps in theoretical and experimental research. Simulating these charged particle interactions is important to a range of scientific disciplines, including biology, chemistry and physics, notably to fusion energy experiments planned for NNSA's National Ignition Facility.

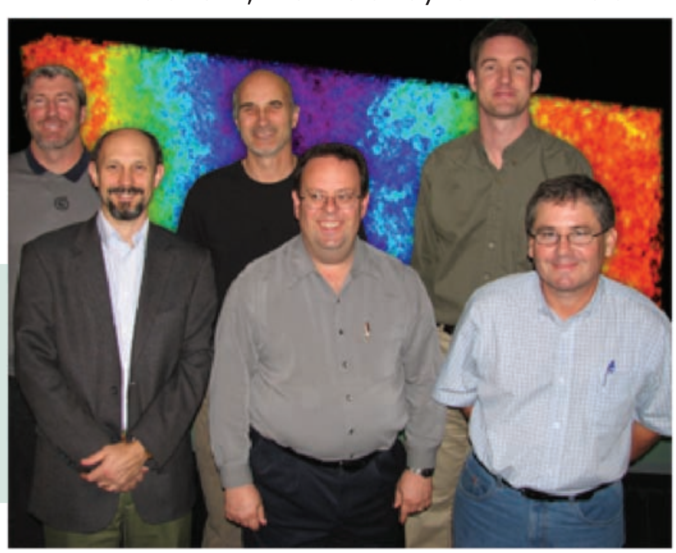

Some members from the team pose in front of the simulation showing the electron potential energy in a 400-million particle molecular dynamics simulation of a hydrogen plasma as it is heated from 500-electron-volts to more than 2-kilo-electron volts by a proton beam.

New Visualization Cluster Supports ASC Data Analysis at Grand Scales

A new visualization cluster at Lawrence Livermore National Laboratory—the 576-node cluster known as Graph—was designed specifically to support interactive data analysis of large, complex datasets generated on the 500-teraFLOPS Sequoia Initial Delivery (ID) System (Dawn) system. Graph provides a memory-rich environment with input/output optimized through dedicated connections to a Lustre file system shared with Dawn. Graph, an Appro system, was delivered to Lawrence Livermore on September 9, 2009.

The Graph visualization cluster installed at Lawrence Livermore.

The Graph cluster nodes each have 24 AMD Opteron processors (quad-socket, six-core Istanbul), for a total of 13,824 processors. To achieve effective data analysis performance, each node has a full complement of 128 gigabytes of RAM for a total of 72 terabytes of memory. Graph successfully combines AMD Opteron processors, InfiniBand interconnect, and the shared ASC TOSS software stack to deliver a stable, reliable, powerful data analysis platform.

Visualization specialists currently manage multi-terabyte datasets with tens of billions of zones, thousands of files per timestep, and hundreds of timesteps. One of the main tools used on Graph thus far has been VisIt, a richly featured visualization and data analysis tool. (See the September 2009 ASC eNews for an article on VisIt performance at scale.) Shortly after Graph was installed, the VisIt team was able to run an 8-trillion-zone problem using 12,000 processors of Graph. This is now the leading example of using VisIt at scale, and it further demonstrates that VisIt and ASC visualization specialists are up to the task of helping scientists analyze datasets far larger than those currently produced (see the VisIt Top 50 list for more information).

Post processing of simulations on ASC Dawn and other LLNL computers will be enabled through Graph's visualization software environment, which supports a wide range of high-performance applications and enables effective remote visualization. The suite of tools includes VisIt, as well as a collection of off-the-shelf tools and a variety of non-commercial, application-specific visualization and data analysis tools. The supporting visualization utilities and infrastructure will include X11, openGL, VTK, movie players, and image viewing and conversion programs, among others. To support the very large datasets being analyzed on Graph, the client server visualization applications is expected to deliver imagery over the network to remote desktops.

Initial B61 Impact Computational Study Completed

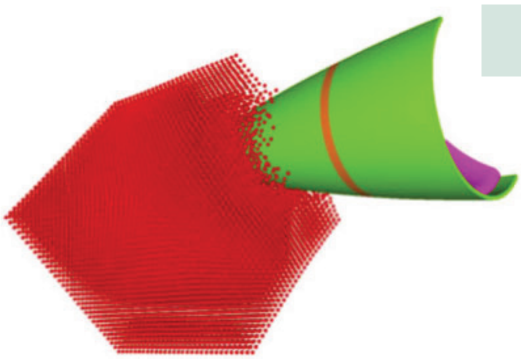

Simulation of radar nose impact into a water target.

An initial computational simulation study has been completed to predict the structural response of a B61 in various impact scenarios (see figure).

This study used Sandia's SIERRA simulation software to predict the structural loads on the B61. New insights were gained using the SIERRA model in a computational sensitivity analysis study to determine which of the B61 material model parameters had the greatest effect on the impact fuze acceleration loads.

This information will be used to guide testing done under Campaign 6 to provide more accurate material property data for use in future SIERRA analyses. Moreover, this SIERRA modeling capability will support current B61 current and future stockpile assessment activities.

Validation Study Provides Increased Confidence

A fire physics computational model validation study has been completed using the thermal/fluid modeling capabilities of Sandia's SIERRA software along with test data from a large (35 meter) hydrocarbon fuel fire. This study provides increased confidence in the ability of SIERRA to model the complex abnormal thermal environments that could be created in a weapon accident scenario.

These SIERRA capabilities will be employed in Sandia's future stockpile assessment activities.

This work also has relevance to other Sandia natural security applications such as large scale hydrocarbon fire scenarios of interest to the Department of Homeland Security.

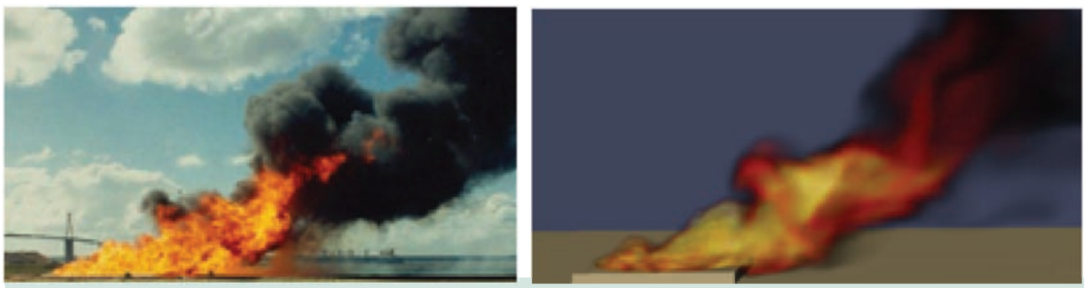

A comparison between the test fire (left) and the SIERRA fire simulation (right). As an example of the validation study results, the test fire had a measured surface emissive power (SEP) of 265 +/- 30 kW/m^2, and the simulation had a predicted SEP of 277 +/- 20 kW/m^2.

Evaluated Nuclear Data Library Supports Current and Future Nuclear Data Needs

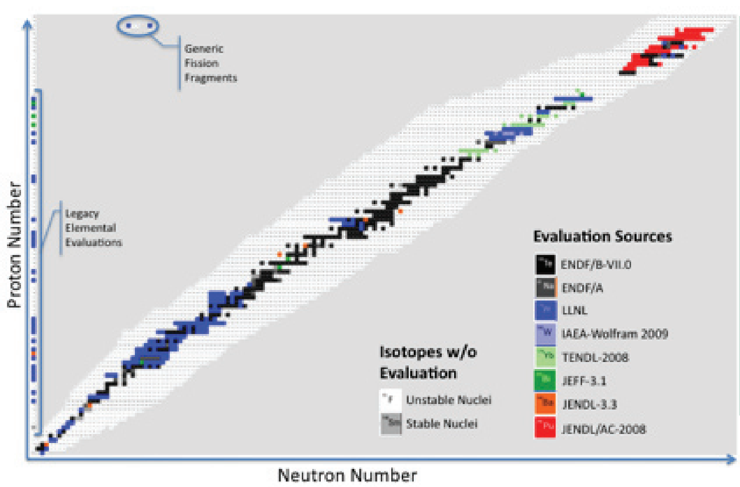

Table of isotopes illustrating the 585 evaluations present as targets for neutron-induced reactions in ENDL2009. Each target isotope is color-coded according to the source library for its evaluation. Isotopes in white are unstable nuclei for which no evaluation exists. Isotopes in light gray are stable nuclei corresponding to the nine stable isotopes for which no evaluation exists. The generic fission fragment and elemental evaluations are maintained for compatibility with archival calculations.

The most complete nuclear database for Monte Carlo and deterministic transport of neutrons and charged particles for use in nuclear reactor, nuclear security, and stockpile stewardship ASC codes was released by Lawrence Livermore National Laboratory (LLNL). The 2009 release of the Evaluated Nuclear Data Library (ENDL2009) supports LLNL's current and future nuclear data needs. This library was assembled with strong support from the ASC Physics Engineering Methods and the Attribution programs, leveraged with support from the Science Campaigns and the DOE/Office of Science's U.S. Nuclear Data Program.

ENDL2009 includes 585 distinct transport-ready evaluations in the neutron sub-library and another 35 evaluations in the charged-particle sub-libraries. It contains many physics improvements for calculating weapon performance, output effects, attribution signatures (part of an ASC Level 2 milestone), key radiochemical diagnostics (part of a Science Campaign Level 2 milestone), and performance of conventional and hybrid fission/fusion reactors. In building this library, the best output from the world's nuclear data efforts were adopted: 46 percent of the library is from the ENDF/B-VII.0 library, 10 percent is from the JENDL libraries, and 8 percent from other libraries. The remaining 36 percent of the neutron sub-library and most of the charged-particle sub-libraries consist of new evaluations developed at LLNL for the ENDL2009 library. In addition, ENDL2009 supports new features such as energy-dependent Q values from fission, average momentum deposition, large-angle Coulomb scattering for all charged particles, support for unresolved resonances, and cross-section covariance data. Finally, this library is LLNL's most highly tested nuclear data release, as LLNL's already rigorous testing regime was strengthened by adding tests against activation ratio measurements and more than 1200 new critical assemblies.

The new library data may be viewed in the Nuclear and Atomic Data System data viewer at http://nuclear.llnl.gov/NADS. A technical report detailing the ENDL2009 release is in preparation.

Los Alamos Authors Recognized as Best Paper Finalist at SC09

A team of authors from Los Alamos National Laboratory and other institutions received recognition for their technical paper "PLFS: A Checkpoint Filesystem for Parallel Applications." The paper by John Bent, et al., was one of three papers selected as best paper finalists out of 261 submissions to the SC09 Technical Papers Program.

Abstract

> Parallel applications running across thousands of processors must protect themselves from inevitable system failures. Many applications insulate themselves from failures by checkpointing. For many applications, checkpointing into a shared single file is most convenient. With such an approach, the size of writes are often small and not aligned with file system boundaries. Unfortunately for these applications, this preferred data layout results in pathologically poor performance from the underlying file system which is optimized for large, aligned writes to non-shared files. To address this fundamental mismatch, we have developed a virtual parallel log structured file system, PLFS. PLFS remaps an application's preferred data layout into one which is optimized for the underlying file system. Through testing on PanFS, Lustre, and GPFS, we have seen that this layer of indirection and reorganization can reduce checkpoint time by an order of magnitude for several important benchmarks and real applications without any application modification.

To read the full paper, go to the ACM Digital Library at http://portal.acm.org/citation.cfm?id=1654059.1654081&dl=ACM&coll=ACM

ASC Booth Wins Award at SC09 in Portland, Oregon

Node exhibit showcasing how NNSA has led the way for US preeminence in HPC. Lower right corner of photo shows the sustainable exhibit award from the SC09 Exhibits committee.

For 14 years, the NNSA ASC Program has had a trade-show booth at the annual international supercomputing conference, SC09, in Portland, Oregon. A tri-lab team — from Lawrence Livermore, Los Alamos, and Sandia national laboratories — created an award-winning booth. In keeping with a new sustainability initiative for the SC09 Exhibits, the team accumulated enough "Green Points" to come in as the third greenest booth at the show that had the largest ever number of industry and research exhibitors in the history of the conference. For example, we gave away reusable water bottles and mouse pads that doubled as scratch pads made from recycled paper. We used electronic exhibits in place of paper and donated the carpet for reuse.

The ASC Program demonstrated the numerous tri-lab advances organized around the booth theme — Advancing High Performance Computing (HPC) for National Security. Exhibits built around the subthemes — energy security, global climate change, technological competitiveness, and national security — highlighted ASC's computing capabilities that advance science, ensure national security, and enhance

Left and Below: Visitors to the NNSA ASC tri-lab booth at SC09 could get information on tri-lab advances in HPC from graphical and electronic displays and hands-on demonstrations.

our nation's economic competitiveness. Modes of communication included video, 3D theatre, demonstrations, small-group discussions, and a node display of ASC machines from as early as the 1997 ASCI Red. US and International and US experts and students in the field of HPC and the press visited the booth. Because pictures are worth a thousand words, see photos at https://computation.llnl.gov/icc/scxy/.

R&D 100 Received for Game-Changing ROSE Compiler

Lawrence Livermore computer scientists have developed a software infrastructure—called a compiler—that promises to speed up the research community's development of new computer software and the science and engineering those applications make possible.

The team of computer scientists received an R&D 100 award for the compiler technology called ROSE at a ceremony in SeaWorld, Florida, on November 12.

ROSE is a user-friendly software infrastructure that allows even average programmers to create their own software analysis, transformation, and optimization tools to speed up the development of new applications and all but eliminate their dependency on compiler experts with specialized knowledge—a requirement that often slowed the development of new software.

Ten years in the making, ROSE is "open source" and is available to users via the Internet free of charge for a little over a year, with a customer base that includes national labs, research universities and industry. The availability of ROSE has enabled improved applications for Livermore research programs across the board from cyber security to nuclear weapons. The development of ROSE was funded by the ASC Program, DOE's Office of Science and the Scientific Discovery through Advanced Computing (SciDAC) program, and Livermore's Laboratory Directed Research and Development.

ROSE is supported by a novel and highly automated testing and release system that permits up to two releases per day. ROSE is released via the SciDAC outreach Web site managed by the DOE Office of Science.

For more about ROSE, see the October edition of Science & Technology Review [https://str.llnl.gov/Oct-Nov09/quinlan.html].

ASC Salutes Matt Hopkins

Like many others in the ASC program, Matt Hopkins is drawn to hard problems. Really hard problems. As the technical lead of a team working to predictively simulate neutron tube (NT) operation at Sandia National Laboratories, he directs the efforts of a highly multi-disciplinary group. NTs are plasma-based devices

On a visit to NASA's original Mission Control at Johnson Space Center in Houston, Matt discussed testing re-entry tiles in their plasma arc jet heater with Martian atmosphere for potential future Mars missions.

to create the correct amount of neutrons at the correct time. His goal is not only to predict NT performance (difficult already), but to develop a framework to understand all of the dependencies in NT performance, including the sensitivity to design and material parameters, and to manufacturing variability. Most modeling and simulation projects target either the manufacturing, or the performance, or the failure, of a system. But since the cause of a failure or the variation of performance can often be attributed to variations in manufacturing, Matt and his team are directly incorporating the manufacturing of NTs as "input" to the performance modeling. "This is a big deal that I don't think people realize, yet. These models are almost always separated, often used in different groups that may not communicate very much. Directly connecting performance or failure to the manufacturing processes that lock in that outcome will be a big step for engineering."

A critical gap in better understanding the performance of NTs and other vacuum arc discharge devices is exactly how, when, where, and why a vacuum arc breaks down. Given their widespread use in all kinds of devices, it seems surprising that little work has been done in simulating vacuum arc breakdown. However, "to make real progress," Matt claims, "you really need to get the right group of folks together. You need physicists, experimentalists, modelers, computational scientists, computer scientists, and HPC experts." Vacuum arc devices begin at vacuum, a very non-continuum environment requiring particle kinetic methods (Particle-In-Cell). After breakdown, there are regions of collisional plasma, requiring other approaches to account for collisions and plasma chemistry (based on Direct Simulation Monte Carlo methods). The system has many different time and spatial scales, and all are important, a sure sign of a "really hard" problem. "I really enjoy inventing new methods, and combining existing ones in new ways, to bridge these different scales."

Matt draws an analogy between the particle collider community and the ASC capability-class simulation work they're doing. As physicists get more energy into colliding particles, they delve deeper into the fundamentals of subatomic physics. Similarly, as his team utilizes more and more powerful HPC systems, they delve deeper into the exact details of vacuum arc breakdown. "It's really exciting. Without the power of these HPC systems we would be much too limited in the size and complexity of system you could simulate for our needs. However, with powerful HPC systems where we can model higher particle counts and include a greater number of computational cells, we are attacking a whole different class of problems (like 3D vacuum arc discharge)."

Matt received his Ph.D. in Mathematics in 1997 from Tulane University. After a short stint as an assistant professor at the Rose-Hulman Institute for Technology, and spending a summer at Sandia as a Visiting Faculty in 1998, he came full time to the Laboratories in 1999. Prior to this current effort, which began in earnest 3 years ago, he was the PI of the SIERRA Aria code module, co-designing an advanced code architecture to support the solution of highly coupled nonlinear systems of PDEs. That work has enabled the consolidation of multiple thermal-fluid simulation methods. Other prior activities include modeling MEMS chemical/biological sensors, non-Newtonian fluids, advanced nonlinear solution methods, and an earlier generation of NT modeling (from which they extrapolated some initial methods for the current work). "I really like working across different disciplines. It's really intriguing to step back and look at all the different work in the abstract and see so many similarities."

Newsletter Points of Contact

Send submittals to:

Denise Sessions—Los Alamos National Laboratory denise@lanl.gov

Andrea Baron—Lawrence Livermore National Laboratory baron1@llnl.gov

Reeta Garber—Sandia National Laboratories ragarbe@sandia.gov

Who's Who

ASC Program Managers—Headquarters

Director, Office of Advanced Simulation and Computing, NA-121.2
Bob Meisner—bob.meisner@nnsa.doe.gov
Thuc Hoang—thuc.hoang@nnsa.doe.gov
Sander Lee—sander.lee@nnsa.doe.gov
Ed Lewis—edgar.lewis@nnsa.doe.gov
Erich Rummel—erich.rummel@nnsa.doe.gov
April Commodore—april.commodore@nnsa.doe.gov
Watti Hill—watti.hill@nnsa.doe.gov

ASC Program Managers—Labs

Los Alamos National Laboratory

John Hopson—jhopson@lanl.gov

Cheryl Wampler—clw@lanl.gov

Lawrence Livermore National Laboratory

Michel McCoy—mccoy2@llnl.gov

Brian Carnes, Deputy Program Manager—carnes1@llnl.gov

Sandia National Laboratories

James Peery—jspeery@sandia.gov

Art Ratzel—acratze@sandia.gov

Developed and maintained by Sandia National Laboratories for NA.121.2

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.